arXiv:1502.05375v1 [cs.DS] 18 Feb 2015

On learning k-parities with and without noise

Arnab Bhattacharyya* Ameet Gadekar† Ninad Rajgopal‡

Department of Computer Science & Automation
Indian Institute of Science

February 19, 2015

Abstract

We first consider the problem of learning k-parities in the on-line mistake-bound model: given a hidden vector $x \in \{0,1\}^n$ with $|x| = k$ and a sequence of "questions" $a_1, a_2, \cdots \in \{0,1\}^n$, where the algorithm must reply to each question with $\langle a_i, x\rangle \pmod 2$, what is the best tradeoff between the number of mistakes made by the algorithm and its time complexity? We improve the previous best result of Buhrman et. al. [BGM10] by an $\exp(k)$ factor in the time complexity.

Second, we consider the problem of learning k-parities in the presence of classification noise of rate $\eta \in (0, 1/2)$. A polynomial time algorithm for this problem (when $\eta > 0$ and $k = \omega(1)$) is a long-standing challenge in learning theory. Grigorescu et al. [GRV11] showed an algorithm running in time $\binom{n}{k/2}^{1+4\eta^2+o(1)}$. Note that this algorithm inherently requires time $\binom{n}{k/2}$ even when the noise rate η is polynomially small. We observe that for sufficiently small noise rate, it is possible to break the $\binom{n}{k/2}$ barrier. In particular, if for some function $f(n) = \omega(1)$ and $\alpha \in [1/2, 1)$, $k = n/f(n)$ and $\eta = o(f(n)^{-\alpha}/\log n)$, then there is an algorithm for the problem with running time $\mathrm{poly}(n) \cdot \binom{n}{k}^{1-\alpha} \cdot e^{-k/4.01}$.

1 Introduction

By now, the "Parity Problem" of Blum, Kalai and Wasserman [BKW03] has acquired widespread notoriety. The question is simple enough to be in our second sentence: in order to learn a hidden vector $x \in \{0,1\}^n$, what is the least number of random examples (a, ℓ) that need to be seen, where a is uniformly chosen from $\{0,1\}^n$ and $\ell = \sum_i a_i x_i \pmod 2$ with probability at least $1 - \eta$? Information-theoretically, x can be recovered only after $O(n)$ examples, even if the noise rate η is close to $1/2$. But if we add the additional constraint that the running time of the learning algorithm be minimized, the barely subexponential running time of [BKW03]'s algorithm, $2^{O(n/\log n)}$, still holds the record of being the fastest known for this problem!

Learning parities with noise is a central problem in theoretical computer science. It has incarnations in several different areas of computer science, including coding theory as the problem of learning random binary linear codes and cryptography as the "learning with errors" problem that underlies lattice-based cryptosystems [Reg09, BV11]. In learning theory, the special case of the problem where the hidden vector x is known to be supported on a set of size k much smaller than n has great relevance. We refer to this problem as *learning k-parity with noise* or k-LPN. Feldman et al. [FGKP09] showed that learning k-juntas, as well as learning 2^k-term DNFs from uniformly random examples and variants of these problems in which the noise is adversarial instead of random, all reduce to the k-LPN problem. For the k-LPN problem, the current record is that of Grigorescu, Reyzin and Vempala [GRV11] who showed a learning algorithm that succeeds with constant probability, takes $\binom{n}{k/2}^{1+(2\eta)^2+o(1)}$ time and uses $\frac{k \log n}{(1-2\eta)^2} \cdot \omega(1)$ samples. When the

*Supported in part by DST Ramanujan Fellowship. Email: arnabb@csa.iisc.ernet.in
†Email: ameet.gadekar@csa.iisc.ernet.in
‡Email: ninad.rajgopal@csa.iisc.ernet.in

noise rate η is close to $1/2$, this running time is improved by an algorithm due to G. Valiant [Val12] that runs in time $n^{0.8k} \cdot \mathrm{poly}(\frac{1}{1-2\eta})$. It is a wide open challenge to find a polynomial time algorithm for k-LPN for growing k or to prove a negative result.

Another outstanding challenge in machine learning is the problem of learning parities without noise in the "attribute-efficient" setting [Blu96]. The algorithm is given access to a source of examples (a, ℓ) where a is chosen uniformly from $\{0,1\}^n$ and $\ell = \sum_i a_i x_i \pmod 2$ with no noise, and the question is to learn x while simultaneously reducing the time complexity and the number of examples drawn by the algorithm. Again, we focus on the case where x has sparsity $k \ll n$. Information-theoretically, of course, $O(k \log n)$ examples should be sufficient, as each linearly independent example reduces the number of consistent k-parities by a factor of 2. But the fastest known algorithm making $O(k \log n)$ samples runs in time $\tilde{O}(\binom{n}{k/2})$ [KS06], and it is open whether there exists a polynomial time algorithm for learning parities that is attribute-efficient, i.e. it makes $\mathrm{poly}(k \log n)$ samples. Buhrman, García-Soriano and Matsliah [BGM10] give the current best tradeoffs between the sample complexity and running time for learning parities in this noiseless setting. Notice that with $O(n)$ samples, it is easy to learn the k-parity in polynomial time using Gaussian elimination.

1.1 Our Results

We first study the noiseless setting. Our main technical result is an improved tradeoff between the sample complexity and runtime for learning parities.

Theorem 1.1. *Let $k, t : \mathbb{N} \to \mathbb{N}$ be two functions[1] satisfying $\log \log n \ll k(n) \ll t(n) \ll n$. For any $\delta > 0$, there is an algorithm that learns the concept class of k-parities on n variables with confidence parameter δ, using $O(kn/t + \log \binom{t}{k} + \log(1/\delta))$ uniformly random examples and $e^{-k/4.01}\binom{t}{k} \cdot \mathrm{poly}(n) \cdot \log(1/\delta)$ running time[2].*

Actually, we prove our result in the mistake-bound model [Lit89] that is stronger than the PAC model discussed above (in fact, strictly stronger assuming the existence of one-way functions [Blu94]). As a consequence, a theorem of the above form also holds when the examples come from an arbitrary distribution. We defer the statement of the result for the mistake-bound model to Section 3.

For comparison, let us quote the closely related result of Buhrman et. al.:

Theorem 1.2 (Theorem 2.1 of [BGM10])**.** *Let $k, t : \mathbb{N} \to \mathbb{N}$ be two functions satisfying $k(n) \leqslant t(n) \leqslant n$. For any $\delta > 0$, there is an algorithm that learns the concept class of k-parities on n variables with confidence parameter δ, using $O(kn/t + \log \binom{t}{k} + \log(1/\delta))$ uniformly random examples and $\binom{t}{k} \cdot \mathrm{poly}(n) \cdot \log(1/\delta)$ running time.*

Thus, in the comparable regime, our Theorem 1.1 improves the runtime complexity of Theorem 3.2 by an $\exp(k)$ factor while its sample complexity remains the same upto constant factors. Note that as t approaches k, our algorithm makes $O(n)$ samples and takes $\mathrm{poly}(n)$ time which is the complexity of the Gaussian elimination approach. On the other hand, if $t = n/\log(n/k)$, our algorithm makes $O(k \log(n/k))$ samples and takes[3] $\exp(-k) \cdot \binom{n/k}{k}$ time (ignoring polynomial factors), compared to the trivial approach which explicitly keeps track of the subset of all the k-weight parities consistent with examples given so far and which makes $O(k \log(n/k))$ samples and takes $O(\binom{n}{k})$ time.

We next examine the noisy setting. Here, our contribution is a simple, general observation that does not seem to have been explicitly made before.

Theorem 1.3. *Given an algorithm $\mathcal{A}$ that learns $\mathrm{PAR}(k)$ over the uniform distribution with confidence parameter δ using $s(\delta)$ samples and running time $t(\delta)$,there is an algorithm $\mathcal{A}'$ that solves the k-LPN problem with noise rate $\eta \in (0, 1/3)$, using $O(s(\delta/2)\log(1/\delta))$ examples and running time $\exp(O(H(3\eta/2) \cdot s(\delta/2) \cdot \log(1/\delta)))) \cdot (t(\delta/2) + s(\delta/2)\log(1/\delta))$ and with confidence parameter δ.*

[1] We assume throughout, as in [BGM10], that k and t are constructible in quadratic time.

[2] The "4.01" can be replaced by any constant more than 4.

[3] By $\exp(\cdot)$, we mean $2^{O(\cdot)}$.

In the above, $H : [0,1] \to [0,1]$ denotes the binary entropy function $H(p) = p \log_2 \frac{1}{p} + (1-p) \log_2 \frac{1}{1-p}$. The main conceptual message carried by Theorem 1.3 is that improving the sample complexity for efficient learning of noiseless parity improves the running time for learning of noisy parity. For instance, if we use Spielman's algorithm as $\mathcal{A}$, reported in [KS06], that learns k-parity using $O(k \log n)$ samples and $O(\binom{n}{k/2})$ running time, we immediately get the following:

Corollary 1.4. *For any $\eta \in (0, 1/3)$ and constant confidence parameter, there is an algorithm for k-LPN with sample complexity $O(k \log n)$ and running time $\binom{n}{k/2}^{1+O(H(1.5\eta))}$.*

For comparison, the current best result of [GRV11] has runtime $\binom{n}{k/2}^{1+4\eta^2+o(1)}$ and sample complexity $\omega(k \log n)$. In the regime under consideration, our algorithm's runtime has a worse exponent but an asymptotically better sample complexity.

The result of [GRV11] requires $\binom{n}{k/2}$ time regardless of how small η is. We show via Theorem 1.3 and Theorem 1.1 that it is possible to break the $\binom{n}{k/2}$ barrier when η is a small enough function of n.

Corollary 1.5. *Suppose $k(n) = n/f(n)$ for some function $f : \mathbb{N} \to \mathbb{N}$ for which $f(n) \ll n/\log\log n$, and suppose $\eta(n) = o(\frac{1}{((f(n))^\alpha \log n)})$ for some $\alpha \in [1/2, 1)$. Then, for constant confidence parameter, there exists an algorithm for k-LPN with noise rate η with running time $e^{-k/4.01+o(k)} \cdot \binom{n}{k}^{1-\alpha} \cdot \mathrm{poly}(n)$ and sample complexity $O(k(f(n))^\alpha)$.*

We note that because of the results of Feldman et. al. [FGKP09], the above results for k-LPN also extend to the setting where the example source adversarially mislabels examples instead of randomly but with the same rate η.

1.2 Our Techniques

We first give an algorithm to learn parities in the noiseless setting in the mistake bound model. We use the same approach as that of [BGM10] (which was itself inspired by [APY09]). The idea is to consider a family $\mathcal{S}$ of subsets of $\{0,1\}^n$ such that the hidden k-sparse vector is contained inside one of the elements of $\mathcal{S}$. We maintain this invariant throughout the algorithm. Now, each time an example comes, it specifies a halfspace H of $\{0,1\}^n$ inside which the hidden vector is lying. So, we can update $\mathcal{S}$ by taking the intersection of each of its elements with H. If we can ensure that the set of points covered by the elements of $\mathcal{S}$ is decreasing by a constant factor at every round, then after $O(\log \sum_{S \in \mathcal{S}} |S|)$ examples, the hidden vector is learned. The runtime is determined by the number of sets in $\mathcal{S}$ times the cost of taking the intersection of each set with a halfspace.

One can think of the argument of Buhrman et al. [BGM10] as essentially initializing $\mathcal{S}$ to be the set of all $\binom{n}{k}$ subspaces spanned by k standard basis vectors. The intersections of these subspaces with a halfspace can be computed efficiently by Gaussian elimination. Our idea is to reduce the number of sets in $\mathcal{S}$. Note that we can afford to make the size of each set in $\mathcal{S}$ larger by some factor C because this only increases the sample complexity by an additive $\log C$. Our approach is (essentially) to take $\mathcal{S}$ to be a random collection of subspaces spanned by αk standard basis vectors, where $\alpha > 1$ is a sufficiently large constant. We show that it is sufficient for the size of $\mathcal{S}$ to be smaller than $\binom{n/\alpha}{k}$ by a factor that is exponential in k, so that the running time is also improved by the same factor. Moreover, the sample complexity increases by only a lower-order additive term.

Our second main contribution is a reduction from noiseless parity learning to noisy parity learning. The algorithm is a simple exhaustive search which guesses the location of the mislabelings, corrects those labels, applies the learner for noiseless parity and then verifies whether the output hypothesis matches the examples by drawing a few more samples. Surprisingly, this seemingly immediate algorithm allows us to devise the first algorithm which has a better running time than $\binom{n}{k/2}$ in the presence of a non-trivial amount of noise. The lesson seems to be that if we hope to beat $\binom{n}{k/2}$ for constant noise rates, we should first address the open question of Blum [Blu96] of devising an attribute-efficient algorithm to learn parity without noise.

2 Preliminaries

Let $\mathrm{PAR}(k)$ be the class of all $f \in \{0,1\}^n$ of Hamming weight k. So, $|\mathrm{PAR}(k)| = \binom{n}{k}$. With each vector $f \in \mathrm{PAR}(k)$, we associate a parity function $f : \{0,1\}^n \to \{0,1\}$ defined by $f(a) = \sum_{i=1}^{n} x_i a_i \pmod 2$.

Let $\mathcal{C}$ be a concept class of Boolean functions on n variables, such as $\mathrm{PAR}(k)$. We discuss two models of learning in this work. One is Littlestone's *online mistake bound* model [Lit89]. Here, learning proceeds in a series of rounds, where in each round, the learner is given an unlabeled boolean example $a \in \{0,1\}^n$ and must predict the value $f(a)$ of an unknown target function $f \in \mathcal{C}$. Once the learner predicts the value of $f(a)$, the true value of $f(a)$ is revealed to the learner by the teacher. The *mistake bound* of a learning algorithm is the worst-case number of mistakes that the algorithm makes over all sequences of examples and all possible target functions $f \in \mathcal{C}$.

The second model of learning we consider is Valiant's famous *PAC model* [Val84] of learning from random examples. Here, for an unknown target function $f \in \mathcal{C}$, the learner has access to a source of examples $(a, f(a))$ where a is chosen independently from a distribution $\mathcal{D}$ on $\{0,1\}^n$. A learning algorithm is said to PAC-learn $\mathcal{C}$ with *sample complexity* s, *running time* t, *approximation parameter* ε and *confidence parameter* δ if for all distributions $\mathcal{D}$ and all target functions $f \in \mathcal{C}$, the algorithm draws at most s samples from the example source, runs for time at most t and outputs a function f^* such that, with a probability at least $1 - \delta$:

$$\Pr_{a \leftarrow \mathcal{D}}[f(a) \neq f^*(a)] < \varepsilon$$

Often in this paper (e.g., all of the Introduction), we consider PAC-learning over the uniform distribution, in which case $\mathcal{D}$ is fixed to be uniform on $\{0,1\}^n$. Notice that for learning $\mathrm{PAR}(k)$ over the uniform distribution, we can take $\varepsilon = \frac{1}{2}$ because any two distinct parities differ on half of $\{0,1\}^n$.

There are standard conversion techniques which can be used to transform any mistake-bound algorithm into a PAC learning algorithm (over arbitrary distributions):

Theorem 2.1 ([Ang88, Hau88, Lit89])**.** *Any algorithm $\mathcal{A}$ that learns a concept class $\mathcal{C}$ in the mistake-bound model with mistake bound m and running time t per round can be converted into an algorithm $\mathcal{A}'$ that PAC-learns $\mathcal{C}$ with sample complexity $O(\frac{1}{\varepsilon}m + \frac{1}{\varepsilon}\log\frac{1}{\delta})$, running time $O(\frac{1}{\varepsilon}mt + \frac{t}{\varepsilon}\log\frac{1}{\delta})$, approximation parameter ε, and confidence parameter δ.*

The k-LPN problem with noise rate η, introduced in Section 1, corresponds to the problem of PAC-learning $\mathrm{PAR}(k)$ under the uniform distribution, when the example source can mislabel examples with a rate $\eta \in (0, 1/2)$. More generally, one can study the k-LPN *problem over* $\mathcal{D}$, an arbitrary distribution. [GRV11] show the following for this problem:

Theorem 2.2 (Theorem 5 of [GRV11])**.** *For any $\varepsilon, \delta, \eta \in (0, 1/2)$, and distribution $\mathcal{D}$ over $\{0,1\}^n$, the k-LPN problem over $\mathcal{D}$ with noise rate η can be solved using $\frac{k\log(n/\delta)\omega(1)}{\varepsilon^2(1-2\eta)^2}$ samples in time $\frac{1}{\varepsilon^2(1-2\eta)^2} \cdot \binom{n}{k/2}^{1+(\frac{\eta}{\varepsilon+\eta-2\varepsilon\eta})^2+o(1)}$, where ε and δ are the approximation and confidence parameters respectively.*

3 In the absence of noise

We state the main result of this section.

Theorem 3.1. *Let $k, t : \mathbb{N} \to \mathbb{N}$ be two functions such that $\log\log n \ll k(n) \ll t(n) \ll n$. Then for every $n \in \mathbb{N}$, there is an algorithm that learns $\mathrm{PAR}(k)$ in the mistake-bound model, with mistake bound at most $(1 + o(1))\frac{kn}{t} + \log\binom{t}{k}$ and running time per round $e^{-k/4.01} \cdot \binom{t}{k} \cdot \tilde{O}\left((kn/t)^2\right)$.*

Using Theorem 1.3, we directly obtain Theorem 1.1. In fact, since Theorem 1.3 produces a PAC-learner over any distribution, a statement of the form of Theorem 1.1 holds for examples obtained from any distribution.

For comparison, we quote the relevant result of [BGM10] in the mistake-bound model.

Theorem 3.2 (Theorem 2.1 of [BGM10]). *Let $k, t : \mathbb{N} \to \mathbb{N}$ be two functions such that $k(n) \leqslant t(n) \leqslant n$. Then for every $n \in \mathbb{N}$, there is a deterministic algorithm that learns* $\mathrm{PAR}(k)$ *in the mistake-bound model, with mistake bound at most $k\lceil \frac{n}{t} \rceil + \log \binom{t}{k}$ and running time per round $\binom{t}{k} \cdot O((kn/t)^2)$.*

Note that their mistake bound is better by a lower-order term which we do not see how to avoid in our setup. This slack is not enough though to recover Theorem 3.1 from Theorem 3.2: dividing t by C roughly multiplies the sample complexity by C and divides the running time by C^k in [BGM10]'s algorithm, whereas in our algorithm, dividing t by C roughly multiplies the sample complexity by C and divides the running time by $(1.28C)^k$.

3.1 The Algorithm

Let $f \in \{0,1\}^n$ be the hidden vector of sparsity k that the learning algorithm is trying to learn. Let $e = \{e_1, e_2, \cdots, e_n\}$ be the set of standard basis of the vector space $\{0,1\}^n$.

Let α be a large constant we set later, and let $T = \alpha t$. Note that $T \ll n$. We define an arbitrary partition $\pi = C_1, C_2, \cdots, C_T$ on the set e into T parts, each of size at most $\lceil n/T \rceil$. Next, let $S_1, \ldots, S_m \subset [T]$ be m random subsets of $[T]$, each of size αk. We choose m to ensure the following:

Claim 3.3. *If $m = \tilde{O}\left(\frac{\binom{T}{\alpha k}}{\binom{T-k}{\alpha k - k}}\right)$, then with nonzero probability, for every set $A \subset [T]$ of size k, $A \subset S_i$ for some $i \in [m]$.*

Proof. This follows from the simple observation that for any fixed $i \in [m]$, $\mathbf{Pr}[A \subset S_i] = \binom{T-k}{\alpha k - k} / \binom{T}{\alpha k}$, and so,

$$\mathbf{Pr}[\exists i \in [m], A \not\subset S_i] = \left(1 - \binom{T-k}{\alpha k - k} \Big/ \binom{T}{\alpha k}\right)^m \leqslant e^{-m\binom{T-k}{\alpha k - k}/\binom{T}{\alpha k}}$$

Choosing $m = 2\frac{\binom{T}{\alpha k}}{\binom{T-k}{\alpha k - k}} \log \binom{T}{k}$ and applying the union bound finishes the proof. □

We fix some choice of $S_1, \ldots, S_m \subset [T]$ that satisfies the conclusion of Claim 3.3 for what follows. In fact, the rest is exactly [BGM10]'s algorithm, which we reproduce for completeness.

For every $i \in [m]$, let $M_i \subset \{0,1\}^n$ be the span of $\bigcup_{j \in S_i} C_j$. Note that $\left|\bigcup_{j \in S_i} C_j\right| \leqslant \alpha k \lceil n/T \rceil \leqslant \alpha k \cdot \left(\frac{n}{T} + 1\right) \leqslant \frac{kn}{t} + \alpha k = (1 + o(1))kn/t$, as $t \ll n$ and α is a constant. So, M_i is a linear subspace containing at most $2^{(1+o(1))kn/t}$ points.

Note that every $f \in \{0,1\}^n$ with $|f| = k$ is contained in some M_i. This is simply because every set of k standard basis vectors is contained in at most k of the T parts in the partition π, and by Claim 3.3, every subset of $[T]$ of size k is contained in some S_i.

Initially, the unknown target vector f can be in any of the M_i's. Consider what happens when the learner sees an example $a \in \{0,1\}^n$ and a label $y \in \{0,1\}$. For $i \in [m]$, let $M_i(a,y) = \{f \in M_i : f(a) = y\}$. $M_i(a,y)$ may be of size 0, $|M_i|$ or $|M_i|/2$. Note that the size of $M_i(a,y)$ can be efficiently found using Gaussian elimination.

We are now ready to describe the algorithm:

- **Initialization:** The learning algorithm begins with a set of affine spaces $N_i, i \in [m]$ represented by a system of linear equations. Initialize the affine spaces $N_i = M_i$ for all $i \in [m]$.
- **On receiving an example $a \in \{0,1\}^n$:** Predict its label $\hat{y} \in \{0,1\}$ such that $\sum_{i \in [m]} |N_i(a, \hat{y})| \geqslant \sum_{i \in [m]} |N_i(a, 1-\hat{y})|$.
- **On receiving the answer from the teacher $y = f(a)$:** Update N_i to $N_i(a,y)$ for each $i \in [m]$.

3.2 Analysis

Before we analyze the algorithm, we first establish a combinatorial claim that is the crux of our improvement:

Lemma 3.4. *If α is a large enough constant,*

$$\frac{\binom{T}{\alpha k}}{\binom{T-k}{\alpha k-k}} \leqslant e^{-k/4.01} \cdot \binom{t}{k}$$

Proof.

$$\begin{aligned}
\frac{1}{\binom{t}{k}} \cdot \frac{\binom{T}{\alpha k}}{\binom{T-k}{\alpha k-k}} &= \prod_{i=0}^{k-1} \frac{k-i}{t-i} \cdot \frac{T-i}{\alpha k-i} \\
&= \prod_{i=0}^{k-1} \frac{\alpha t-i}{\alpha k-i} \cdot \frac{k-i}{t-i} \\
&= \prod_{i=1}^{k-1} \left(1 - \frac{i\left(1-\frac{1}{\alpha}\right)\left(\frac{1}{k-i}-\frac{1}{t-i}\right)}{1+\frac{i}{k-i}\left(1-\frac{1}{\alpha}\right)}\right) \\
&\leqslant \prod_{i=1}^{k-1} \left(1 - \frac{0.999}{1+\frac{k-i}{i}\frac{\alpha}{\alpha-1}}\right)
\end{aligned}$$

where the equalities are routine calculation and the inequality is using that $k(n) \ll t(n)$. Each individual term in the product is strictly less than 1. So, the above is bounded by:

$$\begin{aligned}
&\leqslant \prod_{i=k/(2-\varepsilon)}^{k-1} \left(1 - \frac{0.999}{1+\frac{k-i}{i}\frac{\alpha}{\alpha-1}}\right) \\
&\leqslant \left(1 - \frac{0.999}{1+(1-\varepsilon)\frac{\alpha}{\alpha-1}}\right)^{\frac{1-\varepsilon}{2-\varepsilon}k} \\
&\leqslant \exp\left(-\lg e \cdot \frac{0.999(1-\varepsilon)}{(2-\varepsilon)(1+(1-\varepsilon)\frac{\alpha}{\alpha-1})}k\right) \leqslant e^{-k/4.01}
\end{aligned}$$

for a small enough constant $\varepsilon > 0$ and large enough constant $\alpha > 1$. □

Proof of Theorem 3.1. Fix α to be a constant that makes the conclusion of Lemma 3.4 true.

We first check that the invariant is maintained throughout the algorithm that $f \in \cup_{i\in[m]} N_i$. This holds at initiation by the argument given earlier. After that, obviously, if $f \in N_i$, then $f \in N_i(a, f(a))$ for any $a \in \{0,1\}^n$, and so the invariant holds. Therefore, if the algorithm terminates, it will find the hidden vector f and return it as the solution. The rate of convergence is precisely captured by the number of mistakes learning algorithm makes, which we describe next.

Mistake Bound. Notice that when the algorithm begins, the sum of the sizes of all the affine spaces, $\sum_i |N_i| \leqslant \tilde{O}\left(\frac{\binom{T}{\alpha k}}{\binom{T-k}{\alpha k-k}}\right) 2^{(1+o(1))kn/t}$. Now whenever the learner makes a mistake by predicting $\hat{y} \neq y$, the size of all affine spaces $\sum_i |N_i|$ reduces by a factor of at least 2. This is due to the definition of $\hat{y}$ and the fact that $|N_i(a,\hat{y})| + |N_i(a,1-\hat{y})| = |N_i|$.

Hence, using Lemma 3.4, after at most

$$\log\left(\sum_i |N_i|\right) \leqslant \log\left[\tilde{O}\left(\frac{\binom{T}{\alpha k}}{\binom{T-k}{\alpha k-k}}\right) 2^{(1+o(1))kn/t}\right] \leqslant (1+o(1))kn/t + \log\binom{t}{k} - \Omega(k) + \log O\left(\log\binom{t}{k}\right)$$

mistakes, the size of $\cup_{s\in S} N_s$ will decrease to 1, which by the invariant above will imply that $\cup_{s\in S} N_s = \{f\}$, and hence the learner makes no more mistakes. . Since we assume $k \gg \log\log n$ and $t \ll n$, we can bound the number of mistakes by: $(1+o(1))kn/t + \log\binom{t}{k}$

Running Time. We analyze the running time of the learner for each round. At each round, for a question $a \in \{0,1\}^n$, we need to compute $|N_i(a,0)|$ and $|N_i(a,1)|$ as well as store a representation of the updated N_i. Now, since for each N_i is spanned by at most $\ell = (1+o(1))kn/t$ basis vectors, we can treat each N_i as a linear subspace in $\{0,1\}^\ell$. $N_i(a,0)$ and $N_i(a,1)$ can be computed by performing Gaussian elimination on a system of linear equations involving ℓ variables, which takes $O(\ell^2)$ time. Thus, the total running time is $O(m\ell^2)$, which using Lemma 3.4 is exactly the bound claimed in Theorem 3.1. □

4 In the presence of noise

Recall the k-LPN problem. In this section, we show a reduction from k-LPN to noiseless learning of PAR(k) and its applications.

4.1 The Reduction

We focus on the case when the noise rate η is bounded by a constant less than half.

Theorem 1.3 (recalled) *Given an algorithm $\mathcal{A}$ that learns* PAR*(k) over the uniform distribution with confidence parameter δ using $s(\delta)$ samples and running time $t(\delta)$,there is an algorithm $\mathcal{A}'$ that solves the k-*LPN *problem with noise rate $\eta \in (0, 1/3)$, using $O(s(\delta/2)\log(1/\delta))$ examples and running time $\exp(O(H(3\eta/2)\cdot s(\delta/2)\cdot\log(1/\delta)))\cdot(t(\delta/2)+s(\delta/2)\log(1/\delta))$ and with confidence parameter δ.*

Let $\mathcal{A}(\delta)$ be a PAC-learning algorithm over the uniform distribution for PAR(k) of length n with confidence parameter δ that draws $s(\delta)$ examples and runs in time $t(\delta)$. Below is our algorithm for k-LPN. Here, H denotes the binary entropy function $p \mapsto p\log_2(1/p) + (1-p)\log_2(1/(1-p))$.

NOISY(δ,η)

1: Draw $s' = 20s(\delta/2)\log(1/\delta)$ random examples $(a_1,\ell_1),\ldots,(a_{s'},\ell_{s'}) \in \{0,1\}^n \times \{0,1\}$.
2: **for all** $S \subseteq [s'], |S| \leqslant \frac{3}{2}\eta s'$ **do**
3: **for** $i \in [s']$ **do**
4: **if** $i \in S$ **then** $\tilde{\ell}_i \leftarrow 1-\ell_i$
5: **else** $\tilde{\ell}_i \leftarrow \ell_i$
6: **end if**
7: **end for**
8: $x_S \leftarrow \mathcal{A}(\delta/2)$ applied to examples $(a_1,\tilde{\ell}_1),\ldots,(a_{s'},\tilde{\ell}_{s'})$.
9: **end for**
10: Draw $s'' = 600(s'\cdot H(3\eta/2) + \log(8/\delta))$ random examples $(b_1,m_1),\ldots,(b_{s''},m_{s''}) \in \{0,1\}^n \times \{0,1\}$
11: $S^* \leftarrow \arg\max_{S\subset[s'],|S|\leqslant 3\varepsilon s'/2} |\{i \in [s''] : \langle b_i, x_S\rangle = m_i\}|$
12: **return** x_{S^*}

Proof of Theorem 1.3.

Lemma 4.1. *The sample complexity of* NOISY *is $s' + s'' = O(s(\delta/2)\log(1/\delta))$.*

Proof. Immediate. □

Lemma 4.2. *The running time of* NOISY *is $2^{O(H(3\eta/2)s(\delta/2)\log(1/\delta))}\cdot(t(\delta/2)+s(\delta/2)\log(1/\delta)))$.*

Proof. We use the standard estimate $\sum_{i=0}^{\alpha x}\binom{x}{i} \leqslant 2^{H(\alpha)x}$ for $\alpha \leqslant \frac{1}{2}$. The bound is then immediate. □

Lemma 4.3. *If x is the hidden vector and x^* is output by* NOISY(δ)*, then with probability at least $1-\delta$, $x^* = x$.*

Proof. We make the assumption throughout that $\eta = \Omega(1/s(\delta/2))$, as otherwise, with high probability, the example source won't mislabel any of $s(\delta/2)$ samples and so the reduction is trivial.

Let $T = \{i \in [s'] : \langle a_i, x\rangle \neq \ell_i\}$ be the subset of the s' samples drawn in line 1 that are mislabeled by the example source. By the Chernoff bound: $\mathbf{Pr}[|T| > 3\eta s'/2] \leqslant e^{-\eta s'/12} \leqslant \delta/4$. If $|T| \leqslant 3\eta s'/2$, we have with probability at least $1-\delta/2$, $x_T = x$. Thus, for any $i \in [s'']$, $\mathbf{Pr}_{b_i}[\langle x_T, b_i\rangle \neq m_i] \leqslant \eta$. On the other hand, for all $x_S \neq x_T$, $\mathbf{Pr}_{b_i}[\langle x_S, b_i\rangle \neq \langle x, b_i\rangle] = 1/2$, and so $\mathbf{Pr}_{b_i}[\langle x_S, b_i\rangle \neq m_i] = 1/2$ as the noise is random. Again, using Chernoff bounds,

$$Pr[\exists S \neq T \text{ s.t.} |\{i \in [s''] : \langle b_i, x_S\rangle \neq m_i\}| \leqslant 5s''/12] \leqslant 2^{H(3\eta/2)s'} \cdot e^{-s''/450} < \frac{\delta}{8}$$

On the other hand, for x_T itself, $\mathbf{Pr}[|\{i \in [s''] : \langle b_i, x_T\rangle \neq m_i\}| > 5s''/12] < \frac{\delta}{8}$ by a similar use of Chernoff bounds. So, in all, with probability at least $1-\delta$, x_T will be returned in step 12. □

□

When the noise rate η is more than $1/3$, a similar reduction can be given by adjusting the parameters accordingly. Also, when the distribution is arbitrary but the noise rate is less than $1/4$, a similar reduction can be made to work. In the latter case, $\mathcal{A}$ is invoked with a smaller approximation parameter than the one given to NOISY so that the filtering step in line 10 works.

4.2 Applications

An immediate application of Theorem 1.3 is obtained by letting $\mathcal{A}$ be the current fastest known attribute-efficient algorithm for learning PAR(k), the algorithm due to Spielman[4] [KS06] that makes $O(k \log n)$ samples and takes $O(\binom{n}{k/2})$ time (for constant confidence parameter δ). (We ignore the confidence parameter in this section for simplicity.)

Corollary 1.4 (recalled) *For any $\eta \in (0, 1/3)$ and constant confidence parameter, there is an algorithm for k-*LPN *with sample complexity $O(k \log n)$ and running time $\binom{n}{k/2}^{1+O(H(1.5\eta))}$.*

Proof. Immediate from Theorem 1.3. □

Our next application of Theorem 1.1 uses our improved PAR(k) learning algorithm from Section 3.

Corollary 1.5 (recalled) *Suppose $k(n) = n/f(n)$ for some function $f : \mathbb{N} \to \mathbb{N}$ for which $1 \ll f(n) \ll n/\log\log n$, and suppose $\eta(n) = o(1/((f(n))^\alpha \log n))$ for some $\alpha \in [1/2, 1)$. Then, for constant confidence parameter, there exists an algorithm for k-*LPN *with noise rate η with running time $e^{-k/4.01+o(k)} \cdot \binom{n}{k}^{1-\alpha} \cdot \mathrm{poly}(n)$ and sample complexity $O(k(f(n))^\alpha)$.*

Proof. Let $\mathcal{A}$ be the algorithm of Theorem 1.1 with $t(n) = \lceil n/(f(n))^\alpha \rceil$. The running time of $\mathcal{A}$ is $e^{-k/4.01} \cdot \binom{n}{k}^{1-\alpha} \cdot \mathrm{poly}(n)$ and its sample complexity is $O(k \cdot f(n))^\alpha)$. Now, applying Theorem 1.3, we see that since $H(1.5\eta) = o((f(n))^{-\alpha})$, the running time for NOISY is only a $2^{o(k)}$ factor times the running time of $\mathcal{A}$. This yields our desired result.

□

[4]Though a similar algorithm was also proposed by Hopper and Blum [HB01]

References

[Ang88] Dana Angluin. Queries and concept learning. *Mach. Learn.*, 2(4):319–342, 1988. 4

[APY09] Noga Alon, Rina Panigrahy, and Sergey Yekhanin. Deterministic approximation algorithms for the nearest codeword problem. In Irit Dinur, Klaus Jansen, Joseph Naor, and Jos Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 5687 of *Lecture Notes in Computer Science*, pages 339–351. Springer Berlin Heidelberg, 2009. 3

[BGM10] Harry Buhrman, David García-Soriano, and Arie Matsliah. Learning parities in the mistake-bound model. *Inform. Process. Lett.*, 111(1):16–21, 2010. 1, 2, 3, 4, 5

[BKW03] Avrim Blum, Adam Kalai, and Hal Wasserman. Noise tolerant learning, the parity problem, and the statistical query model. *J. ACM*, 50(4):506–519, 2003. 1

[Blu94] Avrim Blum. Separating distribution-free and mistake-bound learning models over the boolean domain. *SIAM J. on Comput.*, 23(5):990–1000, 1994. 2

[Blu96] Avrim Blum. On-line algorithms in machine learning. In *Workshop on on-line algorithms, Dagstuhl*, pages 305–325. Springer, 1996. 2, 3

[BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *Proc. 52nd Annual IEEE Symposium on Foundations of Computer Science*, pages 97–106, 2011. 1

[FGKP09] Vitaly Feldman, Parikshit Gopalan, Subhash Khot, and Ashok Kumar Ponnuswami. On agnostic learning of parities, monomials, and halfspaces. *SIAM J. on Comput.*, 39(2):606–645, 2009. 1, 3

[GRV11] Elena Grigorescu, Lev Reyzin, and Santosh Vempala. On noise-tolerant learning of sparse parities and related problems. In *Algorithmic Learning Theory*, pages 413–424. Springer, 2011. 1, 3, 4

[Hau88] David Haussler. Space efficient learning algorithms. Technical Report UCSC-CRL-88-2, University of California at Santa Cruz, 1988. 4

[HB01] Nicholas J Hopper and Manuel Blum. Secure human identification protocols. In *Advances in cryptology–ASIACRYPT 2001*, pages 52–66. Springer, 2001. 8

[KS06] Adam R Klivans and Rocco A Servedio. Toward attribute efficient learning of decision lists and parities. *J. Mach. Learn. Res.*, 7:587–602, 2006. 2, 3, 8

[Lit89] Nick Littlestone. From on-line to batch learning. In *Proc. 2nd Annual ACM Workshop on Computational Learning Theory*, pages 269–284, 1989. 2, 4

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. 1

[Val84] Leslie G Valiant. A theory of the learnable. *Comm. Assn. Comp. Mach.*, 27(11):1134–1142, 1984. 4

[Val12] Gregory Valiant. Finding correlations in subquadratic time, with applications to learning parities and juntas. In *Proc. 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 11–20. IEEE, 2012. 2